arXiv:1610.00077v1 [q-bio.OT] 1 Oct 2016

Position and content paradigms in genome rearrangements: the wild and crazy world of permutations in genomics.

Sangeeta Bhatia, Pedro Feijão, Andrew R. Francis

October 4, 2016

Abstract

Modellers of large scale genome rearrangement events, in which segments of DNA are inverted, moved, swapped, or even inserted or deleted, have found a natural syntax in the language of permutations. Despite this, there has been a wide range of modelling choices, assumptions and interpretations that make navigating the literature a significant challenge. Indeed, even authors of papers that use permutations to model genome rearrangement can struggle to interpret each others' work, because of subtle differences in basic assumptions that are often deeply ingrained (and consequently sometimes not even mentioned). In this paper, we describe the different ways in which permutations have been used to model genomes and genome rearrangement events, presenting some features and limitations of each approach, and show how the various models are related. This paper will help researchers navigate the landscape of genome rearrangement models, and make it easier for authors to present clear and consistent models.

1 Introduction

While the order of genes along a chromosome has been used to extract phylogenetic information as far back as 1938 (Dobzhansky and Sturtevant, 1938), it wasn't until almost half a century later that the problem of determining distance between gene arrangements was formalized (Watterson et al., 1982). Right from the beginning, the language and notation of permutations were used, and have been part of the field ever since.

Different arrangements are thought to arise due to various large scale changes in the genome, by which we mean changes affecting more than just a single nucleotide. Examples of such changes include the reversal of a segment of the genome (called an *inversion* or *reversal*), the movement of a segment to a different location on the genome (*transposition*), events that split a chromosome in two or join two chromosomes into one single chromosome (*fission* and *fusion*) and the exchange of segments between different chromosomes (*translocation*). The rearrangement distance is defined to be the minimal number of rearrangement events between a pair of genomes. The set of rearrangement events through which genomes are hypothesised to have changed could consist of a single type of event such as inversion, or a combination of events such as inversion, translocation, fusion and fission.

Representing the contrasting genome arrangements on paper is best done with some form of permutation notation. Rearrangement events, then, are performed on the representation of the arrangement. The action of the rearrangement events depends of course on the choice of notation.

There are two main philosophies of representing genomes as permutations and this has generated some confusion in the literature, as it is frequently not clear which paradigm is being used. In this paper we call these paradigms "position" and "content", and while each generates permutations that can look similar, they have significantly different meanings, and rearrangement events are consequently implemented differently in each. In fact there is a clear relationship between them, and in this article we describe them both, their philosophies, their limitations, and the relationship between them.

The structure of the paper is as follows. We begin in Section 2 with some standard notation for permutations that will be used throughout. The following three sections form the body of the paper and for each paradigm describe respectively how genomes are represented (Section 3); how rearrangements are represented (Section 4); and how the paradigms are related (Section 6).

2 Permutation notation

2.1 Basic definitions

A permutation on a set X is a bijection from X to itself. Without loss of generality, the set X is usually considered to be the set of integers $\mathbf{n} = \{1, 2, \ldots n\}$. A permutation π can be written in "two-line notation" by listing all the elements of the domain and their images under π.

Example 2.1. *Let π be the permutation of the set $\mathbf{n} = \{1, 2, 3, 4, 5\}$ that maps $1 \mapsto 4$, $2 \mapsto 3$, $3 \mapsto 5$, $4 \mapsto 1$, and $5 \mapsto 2$. Its representation in two-line notation is given below:*

$$\pi = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 \\ 4 & 3 & 5 & 1 & 2 \end{pmatrix}.$$

As the first line of the two-line notation is always the same, it is common to use only the bottom row to represent a permutation, for instance writing the permutation π in Example 2.1 as $\pi = [4, 3, 5, 1, 2]$. We will use square brackets and commas to denote this representation of a permutation as a list of images of 1 up to n.

Every permutation can also be expressed as a product of disjoint cycles. We write a cycle in a permutation as $(i_1\ i_2\ \ldots\ i_k)$, meaning i_1 is mapped to i_2, i_2 to i_3, etc, with i_k mapped back to i_1. The cyclic decomposition of the permutation π in Example 2.1 is (1 4)(2 3 5).

The *identity permutation* is the permutation $\iota = [1, 2, \ldots, n]$. Every permutation π has a unique *inverse*, denoted by π^{-1}, that satisfies $\pi^{-1}\pi = \pi\pi^{-1} = \iota$. In cycle notation, the inverse can be obtained by simply writing each cycle down in reverse order. For two-line notation, writing the inverse of π requires mapping each number to its position in the second line of the two-line notation for π.

The set of all permutations on a set of size n forms an algebraic structure called the *symmetric group* usually denoted as S_n. The multiplication operation in the symmetric group is the familiar composition of functions.

It is also possible (and common) to represent genomes in ways that incorporate an orientation on the regions. The set of "signed" permutations is used for this, and this set forms a group called the *hyperoctahedral group*. However there are some other subtleties associated with signed permutations and we defer discussing those to Section 3.1.1.

2.2 Actions

The first complexity in representing permutations in cycle notation comes when one chooses whether the permutation acts on the *left* or on the *right*. While it is common in much of mathematics to think of functions acting on the *left*, so that we might write $f(x)$ for f acting on an element x of the domain, it is also common in the study of group actions to write functions acting on the *right*, so that we might write $(x)f$ or simply xf for f acting on x.

For those unfamiliar with this latter usage, it might begin to make more sense when considering function composition: when acting on the left, fg means we "do g first", while when acting on the right it means "do f first", which accords with the way we read (in "left-to-right" languages). The other intuitive feature of actions on the right for permutations is that the reading of an individual cycle goes from left to right, (so that (2 3 5) means 2 maps to 3, 3 maps to 5, and 5 maps to 2), and this then agrees with the order in which a composition of cycles is read. Note the contrast with the

action on the left: in this case a composition of cycles is read from the right, but within each cycle we read left to right.

Both these conventions appear in the literature on genome rearrangements, and so we will describe both, taking care to make it clear at each point which convention is in play. Note that the map described above for writing a permutation as a product of disjoint cycles is valid regardless of which convention is used, because a) within each cycle we always read left to right, and b) disjoint cycles commute (the action is the same regardless of the order in which they are performed).

The choice of whether our permutations act on the right or on the left will become important when we discuss how to implement rearrangements such as inversions on the genome, using cycle notation. This will be discussed in Section 4.

3 Genome representation

We use the words 'genes' and 'regions' interchangeably to mean 'conserved regions' or what are sometimes called 'synteny blocks' or 'locally colinear blocks' in genome rearrangement literature. Essentially these denote regions of DNA that are present in all of the genomes under study, and whose different arrangements are of interest. The word 'genome' is used to indicate the set of chromosomes involved (and sometimes there is just one).

3.1 The position paradigm – Genomes as maps between positions and regions

The first paradigm we discuss is the explicit use of *position* to describe a chromosome. This approach has been used widely in the literature beginning with the seminal papers by Sankoff et al. (1992); Bafna and Pevzner (1993) to recent developments such as Egri-Nagy et al. (2014b). The description of a chromosome here consists of saying which position a gene or region is in, or equivalently, what region is in each position. These two versions may seem hard to distinguish but they reflect two viewpoints that lead to opposite (or inverse) notation, and so we will spend some time expanding on this theme.

The starting point, adopted widely, is to use the two-line notation described above to put a genome into permutation language. This involves either writing the position numbers along the top of an array, and the region numbers along the bottom, or vice versa:

$$\pi = \begin{pmatrix} 1 & 2 & \cdots & n \\ \pi_1 & \pi_2 & \cdots & \pi_n \end{pmatrix} \quad : \quad \text{either} \quad \begin{pmatrix} \text{positions} \\ \text{regions} \end{pmatrix} \quad \text{or} \quad \begin{pmatrix} \text{regions} \\ \text{positions} \end{pmatrix}$$

If the positions are labelled in sequence around each chromosome, then the bottom row of the "positions to regions" version is just the labels of the regions read along the genome.

This notation, for both representations, effectively treats a chromosome as a map π between the set of positions $\mathbf{n} = \{1, 2, \ldots, n\}$ and the set of regions.

It is common to denote the set of regions also by the integers $1, \ldots n$, which, while natural, can lead to confusion because the map π then looks like a bijection on the set $\mathbf{n}$. It is important to realize that, in this genome representation, this map is not a bijection on the set $\mathbf{n}$, but a one-to-one correspondence between two different sets: one positions, one regions. (Multi-chromosomal genomes are usually modelled as a collection of permutations where each permutation encodes a chromosome (e.g. Kececioglu and Ravi (1995); Hannenhalli and Pevzner (1995))). It is also clear that as permutations, the two representations produce inverses of each other: one maps a position to the region that is in that position, and the other maps a region back to its position.

When we write a permutation (genome arrangement) in cycle form, it will also be different depending on which representation we use.

For example, the genome

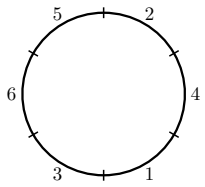

has two-line and cycle notation forms

$$\text{pos} \to \text{reg} : \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 2 & 4 & 1 & 3 & 6 & 5 \end{pmatrix} \;=\; (1\ 2\ 4\ 3)(5\ 6)$$

(the 4-cycle should be read "position 1 has region 2, and position 2 has region 4, and position 4 has region 3" and so on), and

$$\text{reg} \to \text{pos} : \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 3 & 1 & 4 & 2 & 6 & 5 \end{pmatrix} \;=\; (1\ 3\ 4\ 2)(5\ 6)$$

(which should be read "region 1 is in position 3, and region 3 is in position 4", etc). Note that $[(1\ 2\ 4\ 3)(5\ 6)]^{-1} = (1\ 3\ 4\ 2)(5\ 6)$.

Although Watterson et al. (1982) do not explicitly write the gene arrangements in cycle form, their description of a genome consists of mapping regions to positions. Similarly Egri-Nagy et al. (2014b) also write the arrangements with the set of regions constituting the domain and the positions as the co-domain. This choice is more pragmatic from the point of view of software implementation. We will elaborate on this point in Section 4.

In the rest of the literature, genome arrangement as a map from positions to regions is more commonly used (Bafna and Pevzner, 1998; Sankoff et al., 1992; Hannenhalli and Pevzner, 1999). In particular, modeling linear genomes as an ordered list of genes is a natural choice.

Note that this paradigm of genome representations requires referring to an absolute position for each region. This must be chosen *a priori* in order to write down the genome. This choice must then be taken into account when considering the distance between two genomes, as mentioned in Section 4.1.

3.1.1 Incorporating orientation

To incorporate the orientation of DNA into the models, various approaches have been used. The most common is to use *signed permutations*, where the sign of a region represents its orientation. A signed permutation is a permutation of the set $\{-n, \ldots, -1, 1, \ldots, n\}$, where it is common to assume that it satisfies $\pi_{-i} = -\pi_i$ (e.g., Moulton and Steel (2012); Labarre and Cibulka (2011)).

For instance, the genome

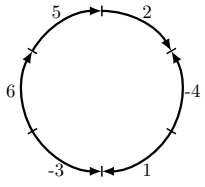

has two-line form

$$\text{pos} \to \text{reg} : \begin{pmatrix} -6 & -5 & -4 & -3 & -2 & -1 & 1 & 2 & 3 & 4 & 5 & 6 \\ -5 & -6 & 3 & -1 & 4 & -2 & 2 & -4 & 1 & -3 & 6 & 5 \end{pmatrix}$$

which in cycle notation form is $(1\ 2\ {-4}\ 3)(-1\ -2\ 4\ -3)(5\ 6)(-5\ -6)$. Note that the cycles come in pairs, and given the constraint that $\pi_{-i} = -\pi_i$, one of each is redundant. Consequently, this

permutation may be abbreviated to (1 2 −4 3)(5 6). The representation in the "reg → pos" form is the inverse of this permutation.

Another way dealing of with orientation is to translate the problem into the realm of unoriented or unsigned permutations. This was done for example by Hannenhalli and Pevzner (1995). Let $\overline{\pi} : \mathbf{n} \to \{-n, \ldots, -1, 1, \ldots n\}$ be a one-to-one function where $|\overline{\pi}(i)| \neq |\overline{\pi}(j)|$ for $i \neq j$, that is, each position i points to a different gene $\overline{\pi}(i)$, and therefore $\overline{\pi}$ represents a signed chromosome. $\overline{\pi}$ can be associated with a permutation π on the set $2\mathbf{n} = \{1, 2, \ldots, 2n\}$ where π is related to $\overline{\pi}$ by Equation (1).

$$\pi := \begin{cases} \pi(2i-1) = 2\overline{\pi}(i) - 1, \quad \pi(2i) = 2\overline{\pi}(i) & \text{if } \overline{\pi}(i) > 0 \\ \pi(2i-1) = 2\overline{\pi}(i), \quad \pi(2i) = 2\overline{\pi}(i) - 1 & \text{if } \overline{\pi}(i) < 0. \end{cases} \tag{1}$$

For instance, the signed chromosome above would be represented by the permutation

$$\pi = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 \\ 3 & 4 & 8 & 7 & 1 & 2 & 6 & 5 & 12 & 11 & 10 & 9 \end{pmatrix}.$$

3.2 The content paradigm – Genomes as maps from regions to regions

The second dominant modeling approach is to view a chromosome as a map from the set of regions to itself.

3.2.1 Genomes as cycles

For example, in what is probably the first formal discussion of reversal distance problem, Watterson et al. (1982) described a chromosome as a set of neighboring gene loci (regions). This view of a chromosome was first formalised in terms of permutations by Meidanis and Dias (2000) who use the cycle structure of a permutation to capture a circular chromosome. The notion of i being "mapped to" j in a cycle is interpreted as region i being followed by region j on the chromosome. Therefore, a circular chromosome is represented by the single cycle $\pi = (\pi_1\ \pi_2\ \ldots\ \pi_n)$. Multi-chromosomal genomes are then represented by many disjoint cycles, one for each chromosome. Meidanis and Dias also used signed permutations to model orientation, representing each chromosome by *two* cycles, one for the direct orientation and other for the reverse complement, with the property that $\pi_{-i} = -\pi_i^{-1}$ (note this is different from the convention for signed permutations described in Section 3.1.1).

For instance, the circular signed chromosome from Section 3.1.1 is modelled by

$$\pi = (2\ {-4}\ 1\ {-3}\ 6\ 5)({-5}\ {-6}\ 3\ {-1}\ 4\ {-2}).$$

This representation has been used to study some rearrangement distances, such as fission, fusion and transposition (Dias and Meidanis, 2001), block-interchange (Huang et al., 2010), and 2-break operations (Feijão and Meidanis, 2013).

Although multi-chromosomal genomes are easier to write down, when compared to the "positional" permutations, it is clear that this formalism is more suitable for dealing with circular genomes than linear genomes, since chromosomes are permutation cycles.

3.2.2 Genomes as adjacencies

In order to develop a permutation model that allows for multi-chromosomal genomes, with both linear and circular chromosomes, it was necessary to use an alternative formulation, where the focus is shifted from the ordering of the genes to the connections between genes. This model became one of the most common ways of representing genomes in the combinatorial community (e.g. Bergeron et al. (2006); Tannier et al. (2009)). A *gene* is defined as an oriented section of the DNA, and its two ends are called its *extremities*. To represent a genome, considered as an arrangement of oriented genes, it is sufficient to note which extremities are adjacent on the genome. An (unordered) pair of extremities that are

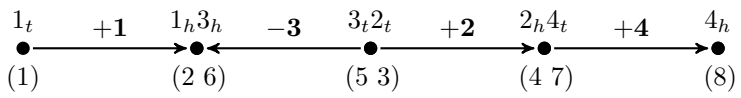

Figure 1: The genome graph of a genome with one linear chromosome containing genes numbered 1, 2, 3 and 4. The vertex set of the graph is the set of adjacencies and extremities, $\{1_t, 1_h3_h, 3_t2_t, 2_h4_t, 4_h\}$. Directed edges are drawn from the tail to the head of the same gene. To express this chromosome as a permutation π, the set of extremities $\{1_t, 1_h, 2_t, 2_h, \ldots, 4_h\}$ is mapped into $\{1, 2, \ldots, 8\}$ (Bhatia et al. (2015)) or into $\{+1, -1, \ldots, +4, -4\}$ (Feijão and Meidanis (2013)). 1_h is connected to 3_h which is captured by the 2-cycle (2 6) in the permutation encoding. The other 2-cycles can be similarly interpreted. The above genome is thus encoded as the permutation (2 6)(5 3)(4 7) or $(-1\ -3)(3\ 2)(-2\ 4)$.

adjacent is referred to as an *adjacency*. An extremity that is not adjacent to any other is the end point of a linear chromosome and is called a *telomere*.

This model can easily uniquely describe a multi-chromosomal genome by its set of adjacencies and telomeres, even if the genome contains both linear and circular chromosomes. The *genome graph* is a graph where the vertices are the adjacencies and telomeres of a genome (sets of one or two gene extremities), and for each gene there is a directed edge from the vertex with the tail of the gene to the vertex with the head of the gene. A *path* in this graph corresponds to a linear chromosome, and a *cycle* to a circular chromosome. It is easy to recover the gene order and orientation in this graph by traversing the components, labelling the edges with the gene label and a plus or minus sign, depending on the direction of the traversal and edge orientation. In the rest of this paper, we refer to this genome representation as the *adjacency list* model.

An algebraic formulation of the adjacency list model was presented by Feijão and Meidanis (2013) and independently by Bhatia et al. (2015). In the algebraic framework, the gene extremities may be represented by the set of signed integers $\{-n, \ldots, -1, 1, \ldots n\}$ (Feijão and Meidanis, 2013) or by mapping the set of $2n$ gene extremeties into the set $\{1, 2, \ldots, 2n\}$ (Bhatia et al., 2015). An adjacency is represented as a 2-cycle and a genome is expressed as a product of disjoint 2-cycles (1-cycles implicitly represent the telomeres and are usually omitted, since they represent fixed points in the permutation).

Figure 1 shows a genome graph consisting of a linear chromosome, together with its permutation representation as an adjacency list.

4 Genome rearrangement events

As we have seen in Section 3, a genome may be viewed as a map between sets of positions and regions, or from a set of regions to itself. In this section we describe how a rearrangement event can be also modelled as a map that acts on a genome, using the language of permutations. The rearrangement event map should therefore be *composable* with the genome, meaning that as functions the appropriate matchings of domains and codomains occurs.

We begin with the position paradigm.

4.1 Rearrangements in the position paradigm

A rearrangement in this paradigm is an action on the *positions* of the genome. So, for instance, an inversion swapping adjacent positions 2 and 3 swaps those two positions regardless of what regions are in those positions. This is really the logic behind this paradigm, as an operation that acted on *regions* 2 and 3 would affect those two regions regardless of where they physically were located on the genome (this is a feature of rearrangements in the "content" paradigm, described in Section 4.2).

Thus, in this view of modeling genomes, a rearrangement event is a permutation on the set of positions.

The observation that the position paradigm involves a map between positions and regions, while a rearrangement is a map from positions to positions, imposes constraints on how they may be composed. These constraints, in turn, depend on whether we wish to write our group actions (rearrangement events) acting on the *left* or on the *right*. The permissible compositions are shown in Table 1.

Table 1: Correct ways to compose a rearrangement such as inversion ("inv") with a genome (permutation). The constraints on the action arise because of the needs for codomain of the first function to match the domain of the second function. For example, after applying a genome permutation representation that takes positions to regions, one cannot then act by a rearrangement function that takes positions to positions.

Permutation representation:	Composition:	Legal action:	Used?
reg $\to$ pos	([reg]) (reg $\to$ pos) $\cdot$ inv	Only on the right	Yes
pos $\to$ reg	(reg $\leftarrow$ pos) $\cdot$ inv ([pos])	Only on the left	Yes
pos $\to$ reg	([pos]) inv $\cdot$ (pos $\to$ reg)	Only on the right	No
reg $\to$ pos	inv $\cdot$ (pos $\leftarrow$ reg) ([reg])	Only on the left	No

Another way to represent the different approaches in the first two rows of Table 1 is given in Figure 2. This figure brings one further subtlety in the difference between the two representations to light. On the left side of the figure, we have "group actions on the right". That is, the group elements multiply (compose) on the right, so that $\rho_1\rho_2$ means that ρ_1 is done first, then ρ_2. It means that if we have a genome π, which is a map from regions to positions, and then we perform several rearrangements on it, first ρ_1, then ρ_2, and so on up to ρ_n, then we can write this:

$$\pi \cdot \rho_1\rho_2 \dots \rho_n.$$

However, on the right hand side of the figure, in which we have "group actions on the left", the same sequence of inversions must be written $\rho_n \dots \rho_1$. This means that performing this sequence of rearrangements on the genome π involves the composition:

$$\pi \cdot \rho_n \dots \rho_2\rho_1.$$

Note that the action of an additional rearrangement ρ_{n+1} with this representation involves inserting it into the *middle* of the above expression so that the new expression is $\pi \cdot \rho_{n+1}\rho_n \dots \rho_2\rho_1$, whereas in the other representation such an additional rearrangement simply involves multiplying on the right: $\pi \cdot \rho_1\rho_2 \dots \rho_n\rho_{n+1}$. For this reason, computational implementations such as Egri-Nagy et al. (2014a) use the former representation as mentioned in Section 3.

A summary of this issue for each paradigm in Table 1 is given in Table 2.

As pointed out in Section 3, Egri-Nagy et al. (2014a) use the regions to positions description of a genome. The choice of the groups acting on the right thus makes this work an example of the first row of Table 1.

On the other hand, in Bafna and Pevzner (1993) and the subsequent papers based on their work (e.g. Caprara (1997)), the genome is described as a map from positions to regions and the rearrangement operators acts on the left. For example, the inversion operator $\rho(i, j)$ that inverts the positions from i to j is the permutation $[1, \dots, i-1, j, j-1, \dots, i+1, i, j+1, \dots]$. Referring to the second row in Table 1, for any position, we first find its image under ρ and then look at the region mapped to this new position under the genome π. For example, ρ maps i to j and π maps j to $\pi(j)$. In the resulting genome, i is mapped to $\pi(j)$. By considering the action of ρ on the other positions, we see that $\rho(i, j)$ acts on a genome $[\pi(1), \dots, \pi(i-1), \pi(i), \pi(i+1), \dots, \pi(j), \pi(j+1), \dots]$ to produce the genome $[\pi(1), \dots, \pi(i-1), \pi(j), \pi(j-1), \dots, \pi(i+1), \pi(i), \pi(j+1), \dots]$.

Table 2: Different ways to represent the genome and the action of rearrangements, with a focus on the composition of rearrangement operators. In the (right hand) Structure column the placement of an additional rearrangement ρ_{n+1} is shown in bold. Note that when the side of the group action and the side the rearrangement acts on the genome clash (one right, one left), new rearrangements must be added into the *middle* of the expression. This is not ideal, and why we recommend use of the structures in the first and last rows, both involving representations of genomes mapping regions to positions.

Structure	Side of		Structure	Example
	group action	rearrangement action		
$(\text{reg} \to \text{pos}) \cdot \text{inv}$	Right	Right	$(x)\pi \cdot \rho_1 \dots \rho_n \boldsymbol{\rho_{n+1}}$	(x is a region)
$(\text{reg} \leftarrow \text{pos}) \cdot \text{inv}$	Left	Right	$\pi \cdot \boldsymbol{\rho_{n+1}} \rho_n \dots \rho_1(x)$	(x is a position)
$\text{inv} \cdot (\text{pos} \to \text{reg})$	Right	Left	$(x)\rho_1 \dots \rho_n \boldsymbol{\rho_{n+1}} \cdot \pi$	(x is a region)
$\text{inv} \cdot (\text{pos} \leftarrow \text{reg})$	Left	Left	$\boldsymbol{\rho_{n+1}} \rho_n \dots \rho_1 \cdot \pi(x)$	(x is a position)

Finally, as mentioned in Section 3.1, the numbering of positions around a circular genome is arbitrary, and it is difficult to talk of two genomes as having the "same" labelling. In fact, the choice of labelling will have a strong effect on the minimal distance, and needs to be taken into account. If there are n regions, then there are n positions and this gives rise to $2n$ distinct labellings of positions on a circle: there are n choices for the first position, and two orientations one could use. Considering these $2n$ alternative labellings is equivalent to considering the action of the *dihedral group* on the set of arrangements, as noted in (Egri-Nagy et al., 2014b; Solomon et al., 2003). Similarly Watterson et al. (1982) also consider the symmetries of a circular genome in determining the minimal distance while Chen and Skiena (1996) factor in the rotational symmetry i.e., the fact that a circular chromosome does not have a fixed 12 o' clock position.

4.2 Rearrangements in the content paradigm

When a genome is modeled as a map from a set of regions to itself, a rearrangement event must be modeled as a similar map and may act on either left or right. The effect of left and right actions will of course be different, but both these formulations have been used in the literature.

Consider for example the model of a circular genome as a cycle (Meidanis and Dias, 2000) where i is mapped to j if region i is adjacent to region j on the genome (Section 3.2.1). Suppose the genome is represented as the k-cycle π, and u is a region on the genome. A transposition event on π that exchanges blocks of length i and j starting at region u, is exchanging the blocks from u to $\pi^{i-1}(u)$ and from $\pi^i(u)$ to $\pi^{i+j-1}(u)$ (with $i+j<k$). To implement this as permutation multiplication, we simply multiply the cycle $\rho = (u\ \pi^i(u)\ \pi^{i+j}(u))$ on the left of π (thinking of the permutations as acting on the left).

For instance, for the genome representation $\pi = (1\ 2\ 3\ 4\ 5\ 6)$, the cycle $\rho = (2\ 4\ 5)$ will transpose regions 2 and 3 with region 4 (recalling here we are acting on the left):

$$\rho\pi = (2\ 4\ 5)(1\ 2\ 3\ 4\ 5\ 6) = (1\ 4\ 2\ 3\ 5\ 6).$$

In this system, the permutation ρ encoding the transposition event depends on the permutation π (the genome) it is acting on. This was commented on by Meidanis and Dias (2000) who note that this makes the problem of determining the rearrangement distance different from the group theoretic problem of determining the word length of a group element under a (fixed) generating set.

4.3 Rearrangement operators on adjacency lists

A popular operator acting on the adjacency list representation of a genome is the "double cut and join" (DCJ) operator introduced by Yancopoulos et al. (2005) and Bergeron et al. (2006). The popularity of this operator can be attributed to its simplicity and the fact that it is able to simulate most of the common rearrangement events that have been observed in comparison of genomes.

Recall the definition of a genome graph from Section 3.2. A DCJ operator acts on a genome graph by choosing two adjacencies at which to "cut" the genome graph. The four cut extremities may now be joined in one of two possible ways, each way resulting in a different rearrangement event on the genome. Other possibilities are: to cut one adjacency and join one of the extremities with a telomere; to cut one adjacency into two telomeres; or the inverse operation of joining two telomeres.

Figure 3 shows a signed reversal of a section being obtained through the double cut and join operator.

Recall from Section 3.2 that a genome can be expressed as a product of disjoint 2-cycles. In this representation, a double cut and join operation takes the form of multiplication by two 2-cycles. Feijão and Meidanis (2013) for example present it as multiplication on the left by two 2-cycles while Bhatia et al. (2015) formulate the double cut and join operation as conjugation by a 2-cycle. For instance, consider a linear chromosome with adjacencies $(i\ j)$ and $(k\ l)$. A DCJ that cuts these adjacencies and joins the adjacencies $(i\ k)$ and $(j\ l)$ can be obtained by conjugating the genome π by the 2-cycle $(k\ j)$, since

$$(j\ k)\big[(i\ j)(k\ l)\big](j\ k) = (i\ k)(j\ l).$$

The same effect can be obtained by left or right multiplication by two 2 cycles,

$$(i\ l)(j\ k)\big[(i\ j)(k\ l)\big] = (i\ k)(j\ l)$$

and equivalently,

$$\big[(i\ j)(k\ l)\big](j\ k)(i\ l) = (i\ k)(j\ l).$$

In fact, there is a simple relationship between these two ways of applying a DCJ operator. In general, when multiplied on the left of a genome π, a DCJ operation is of the format $(u\ v)(\pi u\ \pi v)$ (Feijão and Meidanis, 2013). It is easy to see that $(u\ v)(\pi u\ \pi v)\pi = (u\ v)\pi(u\ v)$, which is the DCJ formulation of Bhatia et al. (2015). The effect of the operator on the genome depends on where the regions u and v lie on the genome.

Feijão and Meidanis (2013) generalise this operator to a k-break operator, where k adjacencies are cut, and k new ones are created from the $2k$ free extremities. As a left operator, the general form of the k-break operator is $(a_1\ a_2\ \ldots\ a_k)(\pi a_k\ \pi a_{k-1}\ \ldots\ \pi a_1)$, which means that a DCJ operation is a *2-break*, the special case for $k = 2$. Interestingly, the k-break operator can also be seen as a conjugation of a genome π by a k-cycle:

$$(a_1\ a_2\ \ldots\ a_k)(\pi a_k\ \pi a_{k-1}\ \ldots\ \pi a_1)\pi = (a_1\ a_2\ \ldots\ a_k)\pi(a_k\ a_{k-1}\ \ldots\ a_1).$$

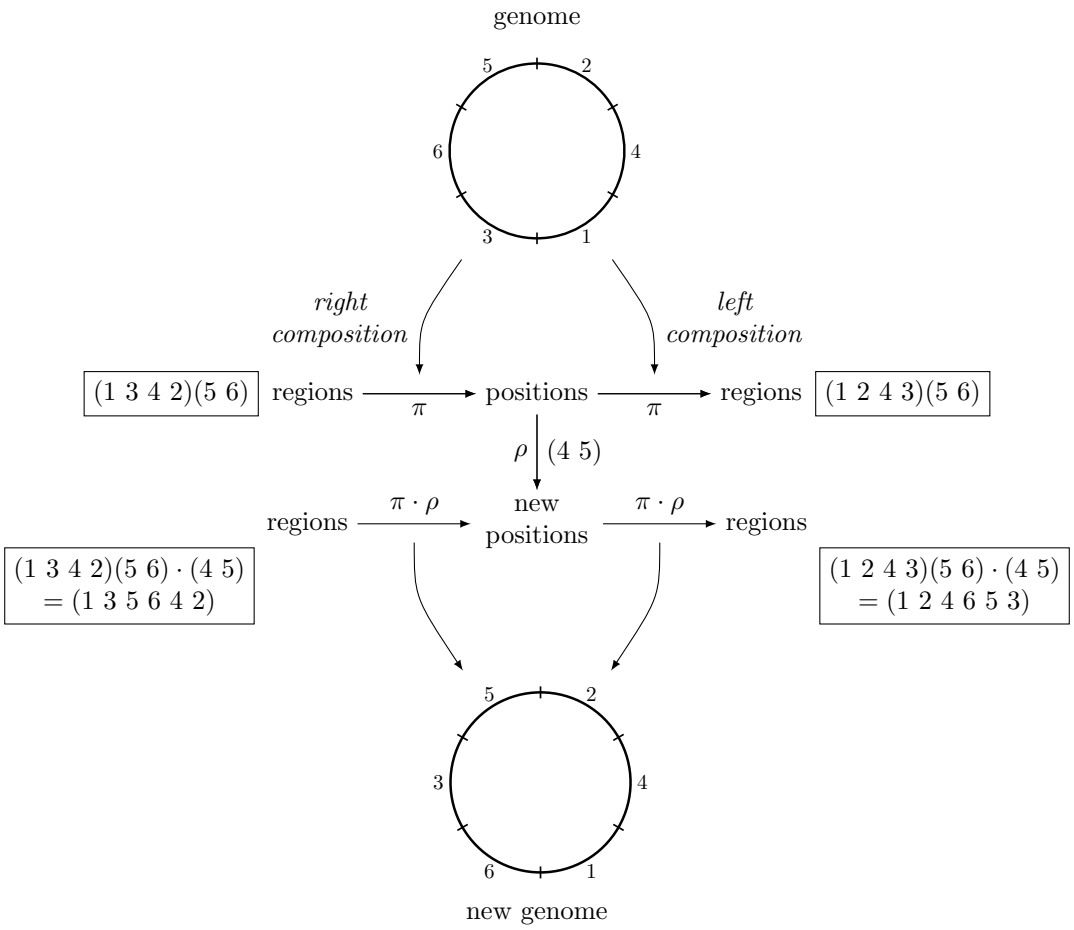

Figure 2: This figure captures the difference in the left and right actions of a rearrangement operator on a circular genome. The left side shows the representation of the genome as a map from regions to positions with a right group action (the first row of Table 1), while the right side shows the genome as a map from positions to regions with a left group action (the second row of the table). The vertical arrow ρ indicates the action of a rearrangement event on the positions of the genome. The effect of this position change on the maps is given in the lower horizontal arrows.

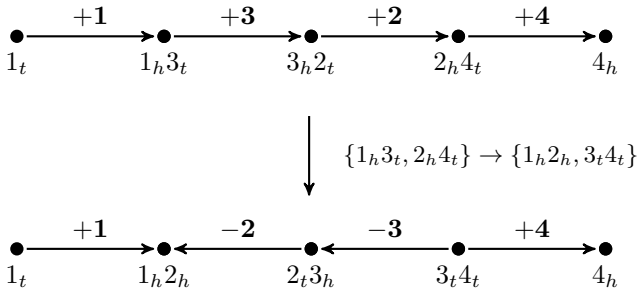

Figure 3: The double cut and join operator can simulate several rearrangement events. In this figure, a double cut is applied at adjacencies 1_h3_t and 2_h4_t, and these extremities are joined to form 1_h2_h and 3_t4_t. In the chromosome at the bottom, the segment between 1_h and 4_t is reversed.

5 The genome rearrangement distance problem

Given two genomes π and σ, we want to find a sequence of rearrangement operations that minimally transforms π into σ. If we consider rearrangements applied on the left, we want to find $\rho_1, \rho_2, \ldots, \rho_k$ such that

$$\rho_k \rho_{k-1} \cdots \rho_2 \rho_1 \pi = \sigma,$$

and k is minimal. The *rearrangement distance* between π and ρ is defined as $d(\pi, \rho) = k$. Finding the minimal k is called the *rearrangement distance problem*, and finding a sequence of rearrangements transforming one genome to the other is commonly called the *rearrangement sorting problem*.

If the rearrangement operators are invertible, then they generate a group. If starting with a single permutation (genome), say the ι permutation, all (unsigned) permutations of the genes can be obtained by the repeated application of the operators under consideration, then the operators generate the symmetric group. Otherwise, the operators generate a subgroup of the symmetric group that is, a group structure that is a subset of the symmetric group.

In the position paradigm of genome representation, the problem of calculating the rearrangement distance is equivalent to finding the *length of a reduced word* for a group element. To see this, recall from above that the distance problem amounts to having two genomes, written π and σ, and wanting to know the minimal number of rearrangements ρ_i such that $\rho_k \rho_{k-1} \cdots \rho_2 \rho_1 \pi = \sigma$. The properties of group multiplication allow us to solve this by considering $s\pi^{-1}$ (now a map from positions to positions), and factorising it in terms of the rearrangement operators. This is precisely the problem of finding a minimal expression for a group element in terms of the generators of the group. Various researchers have noted the connection between the problems of determining rearrangement distance and finding a minimal word for a group element (e.g. see Bafna and Pevzner (1993, 1998)).

6 Going from one genome representation to the other

In this section, we will show how to transform genomes from one representation to another. Unless otherwise noted, the compositions are acting on the left. Clearly, all results could be rewritten with right acting compositions.

6.1 From position to content representation and back

If π is a genome in the regions-to-positions representation (the position paradigm in Section 3.1), then we can transform it into a unique content representation σ by conjugating the n cycle $(1\ 2\ \ldots\ n)$ by π:

$$\pi(1\ 2\ \ldots\ n)\pi^{-1} = \sigma,$$

where σ is the genomes-as-cycles content representation from Section 3.2.1, and where we write our permutations acting on the left (acting on the right swaps π and π^{-1}). This transformation can be understood as follows. Conjugating a k-cycle by an element g of S_n where S_n is the symmetric group on a set of size $\mathbf{n}$, results in applying g to every element of the cycle. That is,

$$g(i_1\ i_2\ \ldots\ i_k)g^{-1} = (g(i_1)\ g_(i_2)\ \ldots\ g(i_k)).$$

Thus, when π conjugates the cycle $(1\ 2\ \ldots\ n)$, we obtain

$$\pi(1\ 2\ \ldots\ n)\pi^{-1} = (\pi(1)\ \pi(2)\ \ldots\ \pi(n))$$

which is exactly the "genomes-as-cycles" content representation.

A similar transformation can be used for signed permutations. If π is a signed permutation in the regions-to-positions representation, satisfying $\pi_{-i} = -\pi_i$ (from Section 3.1.1), then the signed permutation σ in the genomes-as-cycles representation is obtained by

$$\pi(1\ 2\ \ldots\ \ n)(-n\ -n+1\ \ldots\ -1)\pi^{-1} = \sigma.$$

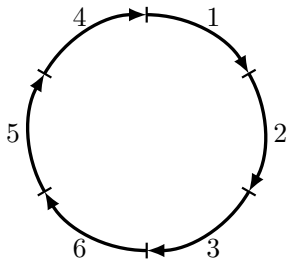

Figure 4: The circular genome encoded by the cycle $\sigma = (1\ 2\ 3\ 6\ 5\ 4)$ in the content paradigm.

$$\begin{pmatrix}1&2&3&4&5&6\\1&2&3&6&5&4\end{pmatrix}\ \begin{pmatrix}1&2&3&4&5&6\\4&1&2&3&6&5\end{pmatrix}\ \begin{pmatrix}1&2&3&4&5&6\\5&4&1&2&3&5\end{pmatrix}$$
$$\begin{pmatrix}1&2&3&4&5&6\\5&6&4&1&2&3\end{pmatrix}\ \begin{pmatrix}1&2&3&4&5&6\\3&6&5&4&1&2\end{pmatrix}\ \begin{pmatrix}1&2&3&4&5&6\\2&3&6&5&4&1\end{pmatrix}$$

Figure 5: The six permutations that encode the genome in Figure 4 as a map from positions to regions, obtained by rotating the labelling of positions. The cyclic (content paradigm) notation for the permutation encoded by these permutations is (1 2 3 6 5 4).

To go from the genomes-as-cycles to the positions-to-regions representation, note that in the genomes-as-cycles representation the information about the position of each region is available relative to the other regions. Since the labeling of the positions can begin at any region, it is easy to see that multiple positions-to-regions permutations will map into the same genomes-as-cycles permutations.

For example, consider the cycle $\sigma = (1\ 2\ 3\ 6\ 5\ 4)$. Viewed as a list of adjacent regions, it encodes the chromosome in Figure 4. Any of the six regions can be thought of as being at position 1, and this assignment then determines the position-to-regions permutation. See Figure 5 for the permutations that correspond to $\sigma = (1\ 2\ 3\ 6\ 5\ 4)$. The first genome π can be obtained from σ by reading σ as the bottom row of a two-line notation (instead of as a cycle), and all the others can be obtained by multiplying π on the right by successive powers of the n-cycle $(n\ n-1\ \ldots\ 2\ 1)$, which corresponds to a "rotation" on the permutation π. There are another n labellings obtained by counting round the circle in the opposite direction (omitted here).

This multiplicity of representations arises from the symmetry inherent in a circular chromosome. Every rotation of a circular arrangement with respect to a fixed reference is equivalent. In fact this symmetry needs to be accounted for while determining rearrangement distance between a pair of circular chromosomes. This was done for example by Watterson et al. (1982) and more recently by Egri-Nagy et al. (2014b) and Serdoz et al. (2016), who also consider the symmetry between a circular permutation and its mirror images.

6.2 From adjacency lists to chromosomes

In the content representation, we saw that there are two interpretations for the cycles of a permutations: the original algebraic formulation of Meidanis and Dias (2000), with cycles corresponding to chromosomes (Section 3.2.1), and the adjacency list representation (Feijão and Meidanis, 2013; Bhatia et al., 2015), with cycles as adjacencies (Section 3.2.2). Although seemingly different, there is a direct relationship between both formulations. Consider the set of signed integers $\{-n, \ldots, -1, 1, \ldots, n\}$,

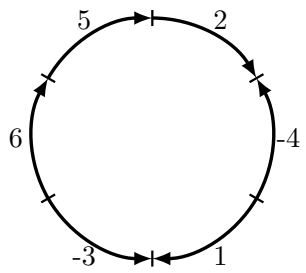

Figure 6: An example genome with oriented regions.

and the permutation $\Gamma = (-1\ 1)(-2\ 2)\cdots(-n\ n)$. Then, going from chromosome to adjacency lists and back is achieved with a right multiplication by Γ (Feijão and Meidanis, 2013) (with permutations acting on the left). For instance, recall that the circular signed chromosome shown in Figure 6 is modelled by $\pi_c = (2\ -4\ 1\ -3\ 6\ 5)(-5\ -6\ 3\ -1\ 4\ -2)$, where each cycle represents a strand of a chromosome, and signs represent gene orientation. Multiplying by Γ on the right we get the adjacency representation π_a:

$$\begin{aligned}\pi_c\Gamma &= (2\ -4\ 1\ -3\ 6\ 5)(-5\ -6\ 3\ -1\ 4\ -2)(-1\ 1)(-2\ 2)\cdots(-6\ 6)\\ &= (-2\ -4)(4\ 1)(-1\ -3)(3\ 6)(-6\ 5)(-5\ 2)\\ &= \pi_a\end{aligned}$$

where negative and (omitted) positive signs represent the head and tail of each gene, respectively.

Since $\Gamma = \Gamma^{-1}$, this transformation works in both directions. It is also possible to apply this transformation for multi-chromosomal genomes, with both linear and circular genomes.

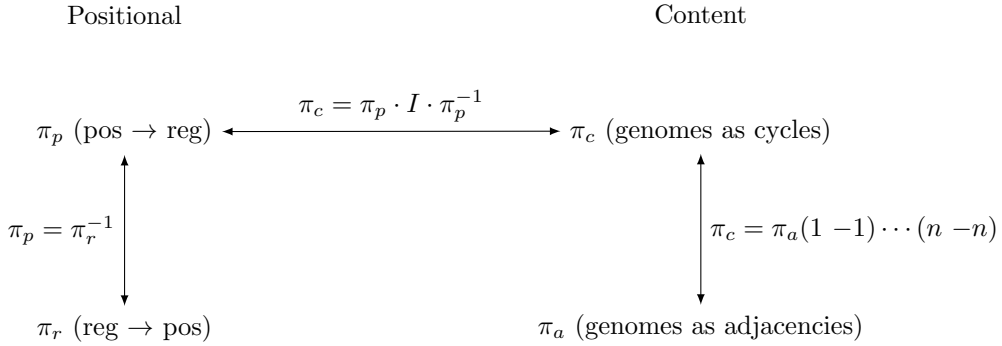

Figure 7: The relationship between permutations arising from different genome representations. $I = (1\ 2\ \ldots\ n)$ for the unsigned case, or $I = (1\ 2\ \ldots\ \ n)(-n\ -n+1\ \ldots\ -1)$ for signed. Note that from π_p to π_c the translation is unique, while the opposite case has multiple possibilities (because there is choice in the labelling of positions). Also, the transformation from π_c to π_a is only possible — or even the modelling of genomes as adjacencies in general — for the signed (oriented) case. All products here assume permutations act on the left.

6.3 Combined approaches

We argue that some approaches in the genome rearrangement literature can, in fact, be interpreted as a combination of the positional and the content paradigms. These approaches are based on interpreting

the cycle graph of a permutation with an algebraic approach.

The *cycle graph* of a permutation π, proposed by Bafna and Pevzner (1998), is a widely used graph where many results on genome rearrangement problems have been obtained. Labarre (2013), extending results from a previous paper (Doignon et al., 2007), proposed a framework that takes a permutation π, representing a genome in the position paradigm, and applies a transformation f, that is an algebraic representation of the cycle graph of π. The transformation $f(\cdot)$ in question is given by

$$f(\pi) = (0\ 1\ \ldots\ n)(\pi_n\ \pi_{n-1}\ \ldots\ \pi_1\ 0)$$

where $\pi_i = \pi(i)$, and a new element 0 is added, to conform with the definition of the cycle graph, where this element is also included. Clearly $f(\pi)$ is always an even permutation, and the authors show that the number of cycles of $f(\pi)$ is the same as the cycle graph of π, which allows them to find new rearrangement distance results based on the number of cycles and on decompositions of $f(\pi)$.

In the content paradigm as defined by Meidanis and Dias (2000), if we want to transform a genome π into another genome σ, we need operations $\rho_1, \rho_2, \ldots, \rho_k$ such that

$$\rho_k \cdots \rho_2 \rho_1 \pi = \sigma \implies \rho_k \cdots \rho_2 \rho_1 = \sigma\pi^{-1}$$

which means that decomposing of the permutation $\sigma\pi^{-1}$ is a way to obtain the rearrangement events, and the number of cycles of $\sigma\pi^{-1}$ is related to many rearrangement distances (e.g., (Dias and Meidanis, 2001; Feijão and Meidanis, 2013; Bhatia et al., 2015)). Now, if we complete the genomes with a 0 element, we have that $\pi = (0\ \pi_1\ \pi_2 \ldots \pi_n)$ and let the target genome σ be the genome where all the elements are "sorted", that is, $\sigma = (0\ 1\ 2\ \ldots\ n)$. Therefore, the permutation $\sigma\pi^{-1}$ is exactly $f(\pi)$ as defined by Labarre. In that sense, we argue that the transformation proposed by Labarre is in fact acting as a translation between the position and content paradigms.

7 Summary and conclusions

In this paper, we have discussed three different models of a genome that are dominant in the genome rearrangement literature. The first of these views a chromosome as a map between a set of positions and a set of regions. The latter two model a genome as a map from a set of regions to itself.

In the position representation of a genome, a rearrangement operator can be thought of as a map from the set of positions to itself. The positions are fixed, as against gene regions which 'move' when a rearrangement operator acts on a genome. Thus a rearrangement operator modeled as a map from positions to positions can have the same form, irrespective of the genome it is acting on. This can simplify the problem of genome rearrangement in some cases, for example if we want to apply a fixed operator such as reversal operator, to a genome.

When a genome is viewed as a map from the set of regions to itself, genome rearrangement operators are also interpreted as maps from the set of regions to itself. This view of a genome has offered great modeling and algorithmic simplicity. For example, contrast the algorithms for sorting via the double cut and join operation Yancopoulos et al. (2005); Bergeron et al. (2006); Feijão and Meidanis (2013); Bhatia et al. (2015) against the sorting by reversals algorithm in Hannenhalli and Pevzner (1999). A disadvantage of this approach is that the action of an operator depends on the structure of the genome on which it is acting. If we wish to fix the operation on the genome, then the set of operations available changes at each step as the genome is being sorted.

References

Vineet Bafna and Pavel A Pevzner. Genome rearrangements and sorting by reversals. In *Proceedings of 1993 IEEE 34th Annual Foundations of Computer Science*, pages 148–157. IEEE, 1993.

Vineet Bafna and Pavel A Pevzner. Sorting by transpositions. *SIAM Journal on Discrete Mathematics*, 11(2):224–240, 1998.

Anne Bergeron, Julia Mixtacki, and Jens Stoye. A Unifying View of Genome Rearrangements. In *Algorithms in Bioinformatics*, pages 163–173. Springer, 2006.

Sangeeta Bhatia, Attila Egri-Nagy, and Andrew R Francis. Algebraic Double Cut and Join. *Journal of Mathematical Biology*, 71(5):1149–1178, 2015.

Alberto Caprara. Sorting by reversals is difficult. In *Proceedings of the First Annual International Conference on Computational Molecular Biology*, pages 75–83. ACM, 1997.

Ting Chen and Steven S Skiena. Sorting with fixed-length reversals. *Discrete Applied Mathematics*, 71(1):269–295, 1996.

Zanoni Dias and J Meidanis. Genome rearrangements distance by fusion, fission, and transposition is easy. In *Proceedings of the 8th International Symposium on String Processing and Information Retrieval (SPIRE2001)*, SPIRE 2001, pages 250–253. Citeseer, 2001.

T Dobzhansky and AH Sturtevant. Inversions in the chromosomes of Drosophila Pseudoobscura. *Genetics*, 23(1):28, 1938.

JP Doignon, Anthony Labarre, and Bd du Triomphe. On Hultman Numbers. *Journal of Integer Sequences*, 10, 2007.

Attila Egri-Nagy, Andrew R Francis, and Volker Gebhardt. Bacterial Genomics and Computational Group Theory: The BioGAP Package for GAP. In *International Congress on Mathematical Software*, pages 67–74. Springer, 2014a.

Attila Egri-Nagy, Volker Gebhardt, Mark M Tanaka, and Andrew R Francis. Group-theoretic models of the inversion process in bacterial genomes. *Journal of Mathematical Biology*, 69(1):243–265, 2014b.

Pedro Feijão and João Meidanis. Extending the Algebraic Formalism for Genome Rearrangements to Include Linear Chromosomes. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 10(4):819–831, 2013.

Sridhar Hannenhalli and Pavel A Pevzner. Transforming men into mice (polynomial algorithm for genomic distance problem). In *Proceedings of 1995 IEEE 36th Annual Foundations of Computer Science*, pages 581–592. IEEE, 1995.

Sridhar Hannenhalli and Pavel A Pevzner. Transforming cabbage into turnip: polynomial algorithm for sorting signed permutations by reversals. *Journal of the ACM (JACM)*, 46(1):1–27, 1999.

Yen-Lin Huang, Cheng-Chen Huang, Chuan Yi Tang, and Chin Lung Lu. An improved algorithm for sorting by block-interchanges based on permutation groups. *Information Processing Letters*, 110 (8-9):345–350, 2010. ISSN 00200190. doi: 10.1016/j.ipl.2010.03.003.

John D Kececioglu and R Ravi. Of mice and men: Algorithms for evolutionary distances between genomes with translocation. In *Symposium on Discrete Algorithms*, volume 604, page 613, 1995.

Anthony Labarre. Lower Bounding Edit Distances between Permutations. *SIAM Journal on Discrete Mathematics*, 27(3):1410–1428, Aug 2013. ISSN 0895-4801. doi: 10.1137/13090897X.

Anthony Labarre and Josef Cibulka. Polynomial-time sortable stacks of burnt pancakes. *Theoretical Computer Science*, 412(8-10):695–702, March 2011. ISSN 03043975. doi: 10.1016/j.tcs.2010.11.004.

João Meidanis and Zanoni Dias. An Alternative Algebraic Formalism for Genome Rearrangements. In *Comparative Genomics*, pages 213–223. Springer, 2000.

Vincent Moulton and Mike Steel. The 'Butterfly effect' in Cayley graphs with applications to genomics. *Journal of Mathematical Biology*, 65(6-7):1267–84, Dec 2012. ISSN 1432-1416. doi: 10.1007/s00285-011-0498-1.

David Sankoff, Guillame Leduc, Natalie Antoine, Bruno Paquin, B Franz Lang, and Robert Cedergren. Gene order comparisons for phylogenetic inference: Evolution of the mitochondrial genome. *Proceedings of the National Academy of Sciences*, 89(14):6575–6579, 1992.

Stuart Serdoz, Attila Egri-Nagy, Jeremy Sumner, Barbara R Holland, Peter Jarvis, Mark M Tanaka, and Andrew R Francis. Maximum likelihood estimates of pairwise rearrangement distances. Preprint, arXiv:1602.03962, 2016.

Andrew Solomon, Paul Sutcliffe, and Raymond Lister. Sorting Circular Permutations by Reversal. In *Workshop on Algorithms and Data Structures*, pages 319–328. Springer, 2003.

Eric Tannier, Chunfang Zheng, and David Sankoff. Multichromosomal median and halving problems under different genomic distances. *BMC Bioinformatics*, 10:120, 2009. ISSN 1471-2105. doi: 10.1186/1471-2105-10-120.

G. A. Watterson, W. J. Ewens, T. E. Hall, and A. Morgan. The Chromosome Inversion Problem. *Journal of Theoretical Biology*, 99(1):1–7, 1982.

Sophia Yancopoulos, Oliver Attie, and Richard Friedberg. Efficient sorting of genomic permutations by translocation, inversion and block interchange. *Bioinformatics*, 21(16):3340–3346, 2005.